CEPREMAP

CENTRE POUR LA RECHERCHE ECONOMIQUE ET SES APPLICATIONS

Document de travail (Docweb) nº 1515

Asset prices and information disclosure under recency-biased learning

Pauline Gandré

Doctoral meetings of the Research in International Economics and Finance (RIEF) network
Cepremap Best Papers Prize 2015

August 2015

Asset prices and information disclosure under recency-biased learning

Pauline Gandré[1]

Abstract: Much of the literature on how to avoid bubbles in international financial markets has addressed the role of monetary policy and macroprudential regulation. This paper focuses on the role of information disclosure, which has recently emerged as a new financial risk management tool. It does so in a consumption-based asset pricing model in which fluctuations in asset prices are persistently driven by time-varying expectations due to learning on the fundamental process from agents who weight more recent observations relative to earlier ones. When the regulator knows the true law of motion driving the fundamental process, perfect information disclosure about the unknown fundamental process straightforwardly rules out non-fundamental fluctuations in asset prices. However, as highlighted by various commentators of the recent financial crisis in 2007-2008, the regulator might also have to learn the true fundamental process and be recency-biased. I investigate the consequences of this assumption on the efficiency of public disclosure about the model actual parameter and identify under which conditions on the regulator learning process, information dissemination could have contributed to significantly reduce the boom and bust episode in the US S&P 500 price index in the run-up to the recent financial crisis. I show that persistent imprecision in the regulator's estimate, which arises as soon as the regulator is recency-biased, can significantly call into question the efficiency of information disclosure for mitigating non-fundamental volatility in asset prices.

Keywords: Asset prices, Bayesian learning, Recency Bias, Information Disclosure, Booms and Busts.

JEL Classification: G15, G12, D83, D84.

Prix des actifs et divulgation d'information en présence de biais d'apprentissage en faveur du présent

Résumé : La littérature étudiant comment éviter les bulles sur les marchés financiers internationaux s'est essentiellement intéressée au rôle de la politique monétaire et de la régulation macro-prudentielle. Ce papier étudie le rôle de la divulgation d'information, qui a émergé récemment comme instrument de gestion des risques financiers. Dans notre modèle, les agents ont une rationalité limitée : leur apprentissage du processus fondamental pondère davantage les observations récentes. Les variations dans les anticipations qui en résultent génèrent des fluctuations persistantes dans le prix des actifs. Un régulateur qui connaît la vraie loi dynamique de l'économie peut éliminer les fluctuations non-fondamentales dans le prix des actifs en diffusant une information non-bruitée. Néanmoins, comme cela a été souligné par de nombreux commentateurs de la crise financière de 2007-2008, le régulateur est lui aussi susceptible de ne pas connaître le vrai processus fondamental et d'être biaisé en faveur du présent. Ce papier identifie sous quelles conditions sur le processus d'apprentissage du régulateur la diffusion d'information aurait pu permettre de réduire significativement la hausse du S&P 500 dans la période qui a précédé la crise financière, suivie de son effondrement. Lorsque le régulateur est biaisé en faveur du présent, la diffusion d'information échoue à réduire les fluctuations non-fondamentales dans les prix des actifs.

Mots-clefs : Prix d'actifs, apprentissage bayésien, biais en faveur du présent, divulgation d'information, Booms et Busts.

---

1 ENS Lyon, université de Lyon, pauline.gandre@ens-lyon.fr, University of Lyon. I am grateful to Julien Azzaz, Camille Cornand, Frédéric Jouneau-Sion, Meguy Kuete, Mirko Wiederholt and Bruno Ziliotto for very helpful discussions at distinct stages of this project, as well as to participants at the 12th Macrofi PhD workshop and the XVth RIEF Doctoral Meetings in International Trade and International Finance. All remaining errors are mine. I also wish to thank the scientific committee of the XVth RIEF Doctoral Meetings for granting me the Best paper prize (shared prize) awarded by CEPREMAP with the support of INFER. This project also benefited from the financial support of the DAAD (German Academic Exchange Service) for funding a research stay at Frankfurt Goethe University.

1 Introduction

During the early 2000s Australian housing and credit bubble, the Reserve Bank of Australia (RBA) implemented original 'open mouth operations' in order to mitigate the steep increase in asset prices and leverage ratios. Such communication policy aimed at warning economic agents against the risk that this increase was not driven by fundamental factors, and thus not sustainable, and might end up in a very costly bust. To this aim, officials from the RBA made public statements highlighting this concern from mid-2002 onwards. This 'public awareness campaign' (Bloxham et al. (2010)), combined with other tools such as monetary policy and regulation tightening, proved rather successful, as the boom ended up in late 2003.

This example suggests that communication policy on the risk of a costly bust when a bubble is identified on a specific financial market –that is when the price deviates from its fundamental value– could be an alternative measure to other standard tools directed at slowing the pace of non-fundamental increase in asset prices. The need for such an alternative measure matters all the more so that much debate subsists on whether monetary policy authorities should consider an asset prices stability objective when setting the interest rate, even after the global financial crisis. Thus, Mishkin (2011) and Woodford (2012), among others, argue that including a financial stability objective in the monetary policy reaction function may generate trade-offs with standard monetary policy objectives, which provides incentives to look for other tools for managing asset prices booms and busts.

In particular, there is room for information dissemination policy as soon as the dynamics of prices are shown to be driven –at least partly– by time-varying expectations. Thus, Williams (2014) emphasizes the role of subjective expectations in explaining fluctuations in asset prices, suggesting that information disclosure aiming at directly impacting expectations and bringing them closer to their rational counterparts could be a relevant alternative tool. Similarly, empirical literature has emphasized that expectations on the future evolution of economic and financial variables are subjective and extrapolated from past realizations of the data. Lovell (1986) shows that the rational expectations assumption does not withstand empirical scrutiny in many economic fields. More specifically, regarding financial markets, de Bondt and Thaler (1985) highlight that stocks of firms which performed badly over the prior years are undervalued whereas stocks of firms which performed well are overvalued. Malmendier and Nagel (2011) provide further evidence that economic agents learn from the past on financial markets.

Drawing upon recent developments in the literature on the impact of learning in macro-finance, this paper provides simple theoretical foundations by relaxing the rational expectations assumption in what regards the law of motion of the exogenous random payoff on a risky asset. This enables to explain several long-standing empirical puzzles in asset pricing theory, unlike the model's rational expectations version, and paves the way for the role of information disclosure about the actual model's parameters in mitigating asset prices fluctuations. Indeed, these fluctuations are shown to be persistently driven by the uncertainty on an unknown parameter of the model, and thus by the endogenous dynamics of beliefs. More specifically, agents learn the location parameter of the law of motion of the risky asset random payoff's growth rate over time, by inferring it from the history of past observations of the economic outcome through Bayesian inference in a standard consumption-based asset pricing model (Lucas (1978)). The only departure from optimal rational behavior that is introduced in the setting, relying on growing empirical evidence (de Bondt and Thaler (1990), Cheung and Friedman (1997), Agarwal et al. (2013), Erev and Haruvy (2013), Gallagher (2014)), is that agents are recency-biased.[1] This means that they react more to recent observations than to earlier ones, as they put recursive weights on past observations, what Fudenberg and Levine (2014) call 'informational discounting'. In particular, Agarwal et al. (2013) evaluate the informational discount rate to be around 90% per month (or equivalently the knowledge depreciation rate is around 10% per month), based on empirical data on the credit card market. In the face of empirical evidence, this assumption of gradual decreasing attention has become rather standard in theoretical literature on financial markets dynamics as well (Bansal and Shaliastovich (2010), Nakov and Nuno (2015)) and enables to generate persistent fluctuations in expectations[2]. In our setting, agents are constrained by limited

[1] Agents are also assumed to be adaptive learners, in the sense that in each period t they condition their expectations on their beliefs in period t but do not take into account the possibility that their future beliefs might change following new realizations of the data (which are still unknown in period t). This assumption is often implicitly made in the Bayesian learning literature and has even more intuitive appeal for recency-biased agents. Indeed, it would be unclear how agents with limited cognitive or technical ability to fully take into account the past would on the contrary be able to account perfectly for all the possible future paths of beliefs when making their forward-looking decision. However, the learning scheme presented here is distinct from stricto sensu adaptive learning as agents are Bayesian learners in the sense that they rely on prior information and take the unknown parameter as a random variable. For a comparison between adaptive and rational Bayesian learning, see Koulovatianos and Wieland (2011).

[2] Initially, recency effects have been studied in models of learning in games in which players choose a best response to what they learned most recently. This characterizes Cournot learning in comparison with fictitious play learning. See for instance Fudenberg and Levine (2014).

cognitive ability to process information and pay less attention to earlier data, what is reflected in recursive discounting of the precision of past observations, in the context of Bayesian inference. In this framework, disclosure of the actual model's parameter by the regulator (may it be a central bank or a financial market regulation authority) can thus bring subjective expectations back to their rational counterparts, causing non-fundamental fluctuations in prices to vanish. Nevertheless, this straightforward mechanism relies on the assumption that rational expectations hold specifically for the regulator. This implies that the regulator knows with certainty the true law of motion of the growth rate of the risky asset's payoff and thus the asset fundamental value. However, the recurrence of very costly unexpected busts in financial markets through economic history – not the least of which being the one which occurred over the period 2007-2009 – tends to prove strikingly that determining an asset fundamental value remains very tricky, even for a regulator having very sophisticated models of risk assessment and big data at his disposal in comparison with standard economic agents. Thus, Andrew Haldane, a Chief Economist with the Bank of England, highlighted that the sophistication of financial risk assessment models in the banking sector in the early 2000s did not help identifying anomalous patterns in financial markets because the regulation authorities suffered from short memory while using these models. Sophisticated stress tests models were only fed with macroeconomic and financial data from the most recent decade, characterized by very low variance, even though the sub-sample distribution was very distinct from the long-run historical distribution (Haldane (2009)). Academic literature has also put recent emphasis on the biases the regulator faces when making its decision (Cooper and Kovacic (2012)). Therefore, it seems justified to investigate the case in which the regulator is not necessarily better informed than economic agents and might not know the true underlying fundamental process of the economy either, neither nor be exempt from recency bias, which may reduce the stabilizing impact of information dissemination. Consequently, this paper aims at assessing to what extent information disclosure about the actual model's parameter can be an efficient tool in helping to mitigate non-fundamental fluctuations in asset prices, depending on the precision of the regulator's estimate.

As an illustration, focusing on one specific period of boom on the stock market and the subsequent bust, the model is calibrated on the US S&P 500 aggregate index and results are compared to monthly data over the 2003-2009 period, that is from the beginning of the boom period on the US stock market to the bust during the global financial crisis, in order to provide quantitative results. It appears that the efficiency of information disclosure for mitigating non-fundamental asset prices fluctuations exhibits strong dependence on the precision of the regulator's estimate of the true parameter (inversely related to its degree of recency bias).[3]

This stems from the fact that when the regulator is recency-biased, even to a really small extent, precision always remains significantly smaller than that in the case where the regulator is not recency-biased, and never reaches infinity. Therefore, this constrains the regulator's public signal on the actual model's parameter to remain imprecise, even if the regulator decides to disclose perfectly its own estimate. This affects the volatility of the price-dividend ratio following public announcement, notably because investors' beliefs react more strongly to new data realizations –whatever the number of past observations– when the regulator's signal is less precise, as investors are then more uncertain about their estimate, even after information disclosure. Counter-factual simulations indeed show that excess volatility in asset prices would be significantly reduced by information dissemination if and only if the regulator's learning process converged to the optimal one, that is the regulator's degree of recency bias tends to zero. This result suggests that information disclosure can only be a relevant tool in helping to mitigate non-fundamental asset prices fluctuations provided that the precision of the regulator's information is high relative to that of standard economic agents, which questions its efficiency.

This paper relates to the growing strand of literature which inserts learning schemes in standard asset pricing models, among which Timmerman (1993), Timmermann (1996), Koulovatianos and Wieland (2011), Adam et al. (2015a) and Adam et al. (2015b). Nevertheless, the learning mechanism is distinct from those investigated in this literature. First, it relies on Bayesian learning (slightly modified in order to account for recency bias), which implies, differently to non-Bayesian learning specifications such as Timmerman (1993), Timmermann (1996) and Adam et al. (2015b) that agents take into account the uncertainty about their estimate of the parameter of the underlying fundamental process when forming their beliefs. Second, differently to Koulovatianos and Wieland (2011) who investigate the impact of learning on the probability of a rare disaster, introducing learning on the mean parameter of the dividend's growth rate allows

[3] Note that differently to usual settings in which the impact of information dissemination on aggregate outcomes was investigated, here informational frictions do not result from imperfect information – meaning that agents observe noisy realizations of the data and implying that information dissemination regards the unknown realized fundamental– but from the relaxation of the rational expectations assumption – meaning that agents do not know at least one parameter of the true laws of motion of economic variables and implying that the information which is disseminated is on the unknown parameters. In this setting, realizations of the fundamental process are perfectly observed by agents.

to generate fluctuations in the price-dividend ratio that are not smooth[4], which is consistent with the high frequency non-monotonous fluctuations observed in the data and allows to replicate additional empirical features. Third, in my model, expectations dynamics are impacted by agents' recency bias. The latter affects the mean value of beliefs, the evolution of the uncertainty associated with these beliefs over time and their degree of persistence. Fourthly, agents derive rationally the equilibrium price as a function of their beliefs on the dividend process, which implies no restriction on the agents' knowledge of the mapping of fundamentals into prices and no additional assumption on the agents' perceived law of motion of prices[5], differently to Adam et al. (2015a) and Nakov and Nuno (2015).

In this setting, a closed-form solution for stock prices can be derived, which enables to make explicit the dependence of the price-dividend ratio on expectations. Furthermore, this closed-form solution has proper microfoundations, as it is directly derived from the representative investor's maximization program.

Even though learning on prices proves very successful in replicating the long-run moments and historical evolution in the US price-dividend ratio thanks to feedback effects (Adam et al. (2015a)), it is less helpful in explaining recent evolution (see also Nakov and Nuno (2015) who replicate well the US price-dividend ratio from 1920 onward but miss the last 25 years). Explaining simultaneously historical features on the US stock market and recent ones with a unique learning scheme and unique parameters seems difficult to achieve. In a sense, this is not very surprising as the underlying fundamental processes are subject to structural breaks and some periods are characterized by much higher uncertainty than others. In particular, strong bust episodes are likely to lead to reassessments of learning processes. Therefore, this paper focuses on the US stock market fluctuations in the run-up to the subprime financial crisis that other models, specified and calibrated in order to target long-term moments, have had more difficulties to explain, suggesting that this period displays specific features. Focusing on some particular episode of boom and bust cycle on the US stock market is of specific relevance when

[4]Indeed, learning only the probability of disaster risk implies that the price-dividend ratio increases monotonously in between two rare disaster realizations. In addition, the authors impose that rare disasters occur in two specific periods whereas I just feed the model with the empirical realizations of the dividend growth rate

[5]In models in which agents learn the stock price process rather than (or in addition to) the dividend process, the perceived law of motion of prices is distinct from the true one in the general case not only regarding the parameters values but also regarding its general specification, and is thus assumed exogenously.

addressing the question of the impact of information disclosure on asset prices movements. Indeed, information dissemination policies, as those implemented by the RBA in the early 2000s, are short-term conjectural policies, responding temporarily to newly identified risks of bubble and of subsequent costly bust on financial markets.

It appears that a parsimonious model with learning on dividends allows to replicate quantitative and qualitative features of the price-dividend ratio evolution in the early 2000s to a rather good extent given the simplicity of the model, even though it generates slightly too much volatility.

In addition, the model allows to replicate (i) the autocorrelation in the price-dividend ratio, (ii) some qualitative features of the dynamics of stock returns and (iii) the positive correlation between expected returns and the current price-dividend ratio in the recent period we focus on, even though slightly overvaluing it.

The impact of several distinct policies on asset prices volatility was investigated in similar frameworks with expectations-driven booms and busts in prices. Thus, Adam et al. (2014) show that a lump-sum tax on financial transactions may deepen volatility in asset prices even though it reduces trading volumes. Winkler (2015) argues that if, under rational expectations, a reaction of monetary policy to asset prices does not improve welfare, it does improve it under subjective expectations.

However, rather paradoxically given that those models explain excess volatility by fluctuations in expectations, much less attention was devoted to the acquisition of information by the regulator through its own –possibly also recency-biased– learning process and to the subsequent role of the regulator's communication policy about the actual fundamental process to private agents. This paper aims at filling this gap. To the best of my knowledge, the impact of the regulator's recency bias on the efficiency of its information disclosure policy about the parameters of the actual law of motion of the economy was never investigated.

Section 2 presents the standard rational expectations model for benchmark purposes and shows how its predictions contradict the data in several ways. Section 3 derives the subjective expectations model with Bayesian learning on the location parameter of the dividend growth rate process and recency bias. Section 4 introduces the regulator's learning process and information disclosure policy about the location parameter and assesses analytically its impact on agents' expectations and aggregate price. Section 5 illustrates the impact of information dissemination depending on the precision of the regulator's signal through counter-factual simulations on the US S&P 500 price-dividend ratio in the run-up to the global financial crisis. Section 6 concludes.

2 The benchmark rational expectations model

I first characterize the model's rational expectations equilibrium. Such an equilibrium can be interpreted as the efficient benchmark equilibrium, arising when the economy is not subject to any kind of friction, in particular informational ones.

The theoretical setting features a simple endowment economy, drawn upon the Lucas (1978)' tree model. In each period, a representative risk-averse agent with CRRA utility function decides what to consume and what to invest in a risky asset (stock) which pays exogenous perishable dividends and a risk-free asset (bonds) which pays an endogenous interest rate. All quantities are expressed in units of a single consumption good. The representative agent's maximization program is the following[6]:

$$\max E_0 \sum_{t=0}^{J} \beta^t \frac{C_t^{1-\gamma}}{1-\gamma} \tag{1}$$

s.t.

$$P_t S_t + C_t + B_t = (P_t + D_t) S_{t-1} + (1 + r_{t-1}) B_{t-1},$$

where β is the discount factor, $\gamma > 0$ is the relative risk aversion coefficient, C_t is consumption in period t, P_t is the stock price, S_t the quantity of stocks, D_t the dividends earned on stocks and r_{t-1} the real interest rate on bonds which is known already at the end of period $t-1$ and B_t the quantity of bonds in period t (the price of bonds is normalized to 1).

[6]As highlighted by Pesaran et al. (2007), in the subjective expectations case with infinitely-lived agents, under general conditions, stock prices do not converge. Pesaran et al. (2007) show that for prices to converge under Bayesian learning, the size of the sample of past observations has to tend towards infinity and grow fast enough relative to the forecast horizon. However, there is no reason for the size of the sample of past observations to be related to the planning horizon. In addition, the previous condition does not allow to study the transition dynamics of the model when the sample of past observations is finite. Therefore, for comparison purposes, I restrict the maximization program under rational expectations to be solved by finitely but very long-lived agents in order to stand close to an infinitely-lived agents set-up. When $t << J$, as is the case in the simulation exercise in Section 5, asset prices are insensitive to the choice of J.

The standard Euler equation with respect to stocks writes:

$$C_t^{-\gamma} = \beta E_t[C_{t+1}^{-\gamma} Z_{t+1}], \tag{2}$$

with $Z_t = \frac{P_t + D_t}{P_{t-1}}$ the gross return on stocks. The Euler equation with respect to bonds writes:

$$C_t^{-\gamma} = \beta(1 + r_t)E_t[C_{t+1}^{-\gamma}]. \tag{3}$$

Stocks are in one unit exogenous supply (as in the Lucas' model), $S_{-1} = 1$ and bonds are in zero net supply. Therefore, the budgetary constraint reduces to $C_t = D_t$. Incorporating this market clearing condition in the above Euler equations characterizes the model's equilibrium. As is standard in the literature, the dividends growth rate follows a log-normal process with parameters d and σ: [7]

$$\log(\frac{D_t}{D_{t-1}}) = d + \sigma\varepsilon_t, \tag{4}$$

with $d > 0$, $\sigma > 0$ and ε_t white noise with unit variance. This implies notably that the persistent component of the dividend growth rate d is constant over time and the process innovations (the transitory component) are unpredictable.

When stock market clears, the first Euler equation (with respect to stocks), reduces to:

$$D_t^{-\gamma} = \beta E_t[D_{t+1}^{-\gamma}(\frac{P_{t+1} + D_{t+1}}{P_t})]. \tag{5}$$

Given the dividend growth process specification, an explicit expression for the price of stocks can be derived. It is as follows:

$$P_t = \delta_t D_t, \tag{6}$$

where $\delta_t = \frac{\beta\theta - (\beta\theta)^{J-t+1}}{1-\beta\theta}$ and $\theta = \exp(d(1-\gamma) + \frac{(1-\gamma)^2\sigma^2}{2})$ (see proof in Appendix A). It appears immediately that the price-dividend ratio in period t (that is δ_t) only depends on the model's fundamental parameters and on time t. Therefore, fluctuations in prices only reflect dividend

[7] See for instance LeRoy and Parke (1992) as an example using the geometric random walk model for dividends on the empirical side and Pesaran et al. (2007) and Adam et al. (2015a) as examples of papers relying on this specification on the theoretical side.

shocks (and time-varying consumption decisions due to the changing distance to the last period). Prices are not impacted by any other variable in the model. Therefore, non-fundamental bubbles cannot arise. This is even more obvious when $t << J$ with J being very high and $\beta\theta < 1$ (which is the case in the calibrated version of the model in section 5), as $\delta_t \to \frac{\beta\theta}{1-\beta\theta}$ and thus the price-dividend ratio tends to be constant. As stated by Lucas (1978), price is then 'a function of the physical state of the economy', being a constant share of current payoffs on stocks. In all cases, there is no role for time-varying expectations, which makes any information disclosure unnecessary. We will see below that this no longer holds when one relaxes the rational expectations assumption.

In addition, under the rational expectations assumption, volatility in stock returns stems only from unpredictable dividend shocks, as:

$$\frac{P_t + D_t}{P_{t-1}} = \frac{1+\delta_t}{\delta_{t-1}} \frac{D_t}{D_{t-1}} = \frac{1}{\beta\theta} \exp(d + \sigma\varepsilon_t). \tag{7}$$

Furthermore, expected returns are constant:

$$E_t[\frac{P_{t+1} + D_{t+1}}{P_t}] = \frac{1}{\beta\theta} \exp(d + 0.5\sigma^2). \tag{8}$$

All these implications of the rational expectations standard asset pricing model are at odds with several features of the data on the US S&P 500, as shown in Appendix B. First, the price-dividend ratio displays large variations and thus excess volatility relative to the prediction of the rational expectations model, suggesting that non-fundamental fluctuations in asset prices do arise. Second, realized stock returns also display excess volatility. Third, expected returns are not constant over time and they are positively correlated with the price-dividend ratio. On the contrary, by allowing agents to learn the location parameter of the logged dividend growth distribution, one can replicate simultaneously all these features. Interestingly, even though alternative theories of excess volatility in asset prices do exist, only the relaxation of the rational expectations assumption enables to generate simultaneously excess volatility and positive correlation between expected returns and the price-dividend ratio, as thoroughly discussed in (Adam et al., 2015a).

3 The subjective expectations model

I now introduce one specific kind of informational friction and assume that agents no longer know the true location parameter of the logged dividend growth process d and learn it over time through recency-biased Bayesian updating.[8]

3.1 Beliefs dynamics

Agents observe the change in dividends, that is the realization of the random variable $y_t = \log(\frac{D_t}{D_{t-1}})$, but they do not observe separately its permanent fixed component d and its transitory component ε_t. Agents are Bayesian learners, which means that they take d as a random variable, they take into account the uncertainty on their estimate. At the beginning of each period, they have prior beliefs on the distribution of d that they update following the new realization of the dividend growth that they observe, according to Bayes' rule. The only departure from rational behavior that I impose is that agents are recency-biased; they have limited ability to pay full attention to earlier data. Less recent data is thus more imprecise for the agents. Therefore, the precision of earlier observations is discounted relative to more recent ones with an informational discount factor $0 \leq \alpha < 1$. Thus, in period t, the precision of observations going back to period $t-k$ is discounted by α^k. α being a discount rate, the higher α, the lower the degree of recency bias.

Bayes' rule applied in period t writes:

$$P(d \mid I_t, \sigma) \propto L(y^t \mid d, \sigma) P(d \mid I_0, \sigma), \tag{9}$$

with $P(d \mid I_t, \sigma)$ the posterior distribution, $P(d \mid I_0, \sigma)$ the prior distribution and $L(y^t \mid d, \sigma)$ the likelihood function. I_t is the information set available at date t, which includes the history of past and current realizations of the logged dividend growth $y^t = \{y_1, y_2, ..., y_{t-1}, y_t\}$. Recency bias modifies the Bayesian inference process only to the extent that the precision of past realizations is discounted. Thus, only parameters are affected.

[8] The location parameter of the logged dividend growth is the only parameter agents have to learn. They know that dividend growth follows a log-normal distribution and they know the dispersion parameter. Assuming that agents know the true precision parameter allows to identify directly the impact of the evolution in one specific belief on asset prices as agents only learn one parameter. In addition, with standard conjugate priors, the posterior distribution for the variance does not exist as shown in Pesaran et al. (2007).

Given that the logged dividend growth process follows a normal distribution, a natural prior distribution for d is the normal conjugate prior, which allows to derive the posterior distribution analytically as it has the same form than the prior distribution. The prior distribution thus writes $P_0 \sim N(m_0, \sigma_0)$. Therefore, applying Bayes' rule under recency-biased learning yields the posterior distribution at the end of period t, $d \backsim N(m_t, \sigma_t)$ with:

$$m_t = \frac{y_t * \rho + y_{t-1}\alpha\rho + y_{t-2}\alpha^2\rho + ... + m_0 * \tau_0 * \alpha^t}{(1 + \alpha + \alpha^2 + ... + \alpha^{t-1}) * \rho + \alpha^t * \tau_0}, \tag{10}$$

with $\tau_0 = \frac{1}{\sigma_0^2}$ the prior precision. The posterior mean belief is the average of the discounted prior belief and of past and current observations weighted by their respective discounted precision in time t (the precision associated with each observation or the prior evolves in each period because of the recency-bias).

$$\sigma_t^2 = \frac{1}{(1 + \alpha + \alpha^2 + ... + \alpha^{t-1}) * \rho + \alpha^t * \tau_0}. \tag{11}$$

See Proof in Appendix C.

As $\alpha < 1$, this reduces to $\sigma_t^2 = \frac{1}{\frac{\rho(1-\alpha^t)}{(1-\alpha)} + \alpha^t \tau_0}$. This term represents the uncertainty on the posterior estimate of d in time t. Unsurprisingly, it decreases in the informational discount rate α (which means that it increases in the degree of recency bias). This is related to the fact that when the weight allocated to earlier data is higher, the precision of the information on the true parameter provided by each past realization of the data is higher, and thus the overall precision is higher.

As is standard in learning models, I now assess how such subjective beliefs nest inside rational expectations.

Proposition 2: Subjective expectations converge to rational ones in the limit if and only if $\alpha = 1$.

Indeed,

$$\lim_{t \to +\infty} \sigma_t^2 = \frac{1}{(1 + \alpha + \alpha^2 + ... + \alpha^{t-1}) * \rho + \alpha^t \tau_0} = 0 \Longleftrightarrow \alpha = 1. \tag{12}$$

Under standard Bayesian learning ($\alpha = 1$), uncertainty on the posterior estimate of d in period t is equal to $\frac{1}{t * \rho + \tau_0}$. Therefore, when the sample of observations increases, uncertainty on the estimate of d decreases and converges to zero. Thus, in the limit, agents are no longer uncertain regarding their mean belief and the law of large numbers ensures that the Bayesian estimate of

d is equal to $\lim_{t\to+\infty} m_t = d$. The posterior estimate of d then converges to its true value when the number of past observations tends towards infinity.

Under recency-biased Bayesian learning ($\alpha \neq 1$), beliefs never converge to rational expectations ones. Indeed, $\lim_{t\to\infty} \frac{1}{\frac{\rho(1-\alpha^t)}{(1-\alpha)}+\alpha^t\tau_0} = \frac{1-\alpha}{\rho}$. Thus, precision on the estimate is never infinite and agents' estimates continue to evolve following new realizations of logged dividend growth even in the limit.

Time-varying subjective expectations then impact the price-dividend ratio, as shown in what follows.

3.2 Closed-form solution for stock prices under recency-biased Bayesian learning

The agents' maximization program now writes

$$\max E_0^P \sum_{t=0}^{J} \beta^t \frac{C_t^{1-\gamma}}{1-\gamma} \tag{13}$$

s.t.

$$P_t S_t + C_t + B_t = (P_t + D_t)S_{t-1} + (1 + r_{t-1})B_{t-1},$$

with P the representative agent's subjective probability measure.

The first-order condition with respect to stocks now depends on this subjective probability measure and writes: $C_t^{-\gamma} = \beta E_t^P[C_{t+1}^{-\gamma} \frac{P_{t+1}+D_{t+1}}{P_t}]$. In combination with the market clearing condition on the stock market, this yields the subjective expectations equilibrium stock price.

Under recency-biased Bayesian learning, stock prices now write:

$$P_t = D_t \sum_{j=1}^{J-t} \beta^j E_t^P[(\frac{D_{t+j}}{D_t})^{1-\gamma}], \tag{14}$$

(see details in Appendix D). Hence, conditionally on d and σ and information up to period t, we have:

$$E^P[(\frac{D_{t+j}}{D_t})^{1-\gamma} \mid I_t, d, \sigma] = \exp[(1-\gamma)jd + 0.5(1-\gamma)^2 j\sigma^2]. \tag{15}$$

Therefore, when d is no longer known, agents integrate the previous expression over the whole distribution of d:

$$E^P[\exp((1-\gamma)jd + 0.5(1-\gamma)^2\sigma^2 j^2) \mid I_t, \sigma] = \exp(0.5(1-\gamma)^2\sigma^2 j)E[\exp((1-\gamma)jd) \mid I_t, \sigma]. \quad (16)$$

As d is believed to follow a normal distribution with parameters m_t and σ_t in period t, $\exp((1-\gamma)jd) \sim Log - N((1-\gamma)jm_t, (1-\gamma)j\sigma_t)$. Therefore,

$$E^P[(\frac{D_{t+j}}{D_t})^{1-\gamma} \mid I_t, \sigma] = \exp(0.5(1-\gamma)^2\sigma^2 j^2)\exp((1-\gamma)jm_t + 0.5(1-\gamma)^2 j^2\sigma_t^2). \quad (17)$$

Eventually, this yields the following pricing function.

Proposition 3 Under recency-biased learning, stock prices write:

$$P_t = D_t \sum_{j=1}^{J-t} \beta^j \exp((1-\gamma)m_t j + 0.5(1-\gamma)^2 j(\sigma_t^2 j + \sigma^2)).\text{[9]} \quad (18)$$

It is immediate to see that under learning, the price-dividend ratio is no longer a constant share of dividends when $t << J$, nor a function only of the model's parameters and of the time period in the general case. Indeed, it now depends on the hyperparameters m_t and σ_t which are updated every period and are therefore time-varying. Therefore, fluctuations in the price-dividend ratio over time –that is, fluctuations in prices which do not reflect directly changes in dividend realizations– are driven by expectations. This generates excess volatility in asset prices. Comparative statics show that the price-dividend ratio in t increases in m_t if and if only $\gamma < 1$, increases in σ_t, in σ, in β and increases in γ if $m_t < 0$ (see Appendix E).

In addition, under subjective expectations, stock returns are likely to display greater volatility than under rational expectations. Indeed,

$$\frac{P_t + D_t}{P_{t-1}} = \frac{1+\delta_{t,SE}}{\delta_{t-1,SE}}\frac{D_t}{D_{t-1}} = \frac{1+\delta_{t,SE}}{\delta_{t-1,SE}}\exp(d + \sigma\varepsilon_t), \quad (19)$$

[9] It is obvious that when $j \to \infty$, the term of the sum depending on j does not converge; that is why J is assumed to be finite.

with $\delta_{t,SE} = \sum_{j=1}^{J-t} \beta^j \exp((1-\gamma)m_t j + 0.5(1-\gamma)^2 j(\sigma_t^2 j + \sigma^2))$. Differently to the rational expectations case, volatility in stock returns does not only stem from volatility in the dividends growth rate, as the ratio $\frac{1+\delta_{t,SE}}{\delta_{t-1,SE}}$ is now time-varying and reflects the additional impact of dividend shocks on expectations.

Regarding expected returns, they are now given by:

$$E_t[\frac{P_{t+1}+D_{t+1}}{P_t}] = \frac{1}{\delta_{t,SE}}(\sum_{j=1}^{J-t-1}[\beta^j \exp(m_t((1-\gamma)j+1) + 0.5((1-\gamma)j+1)^2\sigma_t^2 \\ +0.5((1-\gamma)^2 j+1)\sigma^2)] + \exp(m_t + 0.5(\sigma_t^2+\sigma^2))).$$

(see Proof in Appendix F). Current prices (in the denominator) are positively correlated with the current price-dividend ratio, as a higher price-dividend ratio reflects more optimistic beliefs on the dividend process and thus leads to higher demand for stocks. Similarly, expected future prices and dividend payments (in the numerator) are positively correlated with more optimistic beliefs, and through this, they are also positively correlated with the current price-dividend ratio. This generates the possibility for a correlation between expected returns and the price-dividend ratio – even though the sign of this correlation is analytically ambiguous–, as evidenced in survey data.

Recency-biased Bayesian learning thus allows to match several qualitative features of the data which are not replicated by the rational expectations benchmark model. The learning process implies that non-fundamental fluctuations in asset prices arise as soon as uncertainty on the estimated parameters is not null, bringing the equilibrium price-dividend ratio away from the value it takes in the efficient rational expectations model. Such fluctuations bring the stock price away from its fundamental value and generates additional volatility relative to that in dividends, due to informational frictions. To this respect, those fluctuations are inefficient. As they are driven by expectations, a natural policy to mitigate these fluctuations is to bring expectations closer to their rational expectations counterpart through communication policy and thus information disclosure.[10] In the next section, I investigate the impact of information disclosure from a regulator (a central bank or a financial market authority) which does not know the true parameter of the

[10] Due to some simplifying assumptions such as a representative investor and exogenous stock supply which enable to obtain closed form solutions easily interpretable, proper welfare analysis cannot be performed in our simple set-up. However, due to their non-fundamental dimension, it matters to assess how to mitigate those inefficient fluctuations.

logged dividend growth process either and learns it through Bayesian updating as well.

4 Information disclosure and asset prices

I now model the regulator's own inference process on the true location parameter of the logged dividend growth.

4.1 The regulator's estimate

It learns it through the same inference process as economic agents, except that its degree of recency bias $0 \leq \alpha_R \leq 1$ and its prior distribution $d \backsim N(m_{R,0}, \sigma_{R,0})$ are not restricted to be similar to that of participants in the stock market. In the limiting case in which $\alpha_R = 1$, the regulator does not suffer from recency bias. Therefore, in that case, even if the regulator is not omniscient and does not know the true parameter, its estimate converges to the true one in the limit.

The regulator's posterior distribution for d is $d \sim N(m_{R,t}, \sigma_{R,t})$ with:

$$m_{R,t} = \frac{y_t \rho + y_{t-1} \alpha_R \rho + ... + y_1 \alpha_R^{t-1} \rho + m_{R,0} \alpha_R^t \tau_{R,0}}{(1 + \alpha_R + ... + \alpha_R^{t-1}) \rho + \alpha_R^t \tau_{R,0}}, \tag{20}$$

where $\tau_{R,0} = \frac{1}{\sigma^2_{R,0}}$ and

$$\sigma^2_{R,t} = \frac{1}{(1 + \alpha_R + ... + \alpha_R^{t-1}) * \rho + \alpha_R^t \tau_{R,0}}. \tag{21}$$

When $\alpha_R = 1$, the regulator is not recency-biased:

$$m_{R,t} = \frac{(y_t + y_{t-1} + ... + y_1) \rho + m_0 \tau_{R,0}}{t * \rho + \tau_{R,0}}, \tag{22}$$

and

$$\sigma^2_{R,t} = \frac{1}{t * \rho + \tau_{R,0}}. \tag{23}$$

We can now investigate the impact of information disclosure about the actual parameter from the regulator on agents' expectations and stock prices.

4.2 Information disclosure, agents' expectations and volatility: some analytical results

The regulator's objective is to minimize non-fundamental volatility in asset prices, that is, to minimize the volatility of the price-dividend ratio, which depends on agents' expectations and can thus be impacted by affecting expectations. To achieve this aim, the regulator can disclose information to agents on the distribution of the true model's parameter $d \sim N(m_t(\alpha_D), \sigma_t(\alpha_D))$ with:

$$m_t(\alpha_D) = \frac{y_t\rho + y_{t-1}\alpha_D\rho + y_{t-2}\alpha_D^2\rho + ... + m_{R,0}\alpha_D^t\tau_{R,0}}{(1 + \alpha_D + \alpha_D^2 + ... + \alpha_D^{t-1})\rho + \alpha_D^t\tau_{R,0}},$$

and

$$\sigma_t(\alpha_D)^2 = \frac{1}{(1 + \alpha_D + ... + \alpha_D^{t-1}) * \rho + \alpha_D^t\tau_{R,0}}.$$

α_D is the informational discount factor used to derive the parameters disclosed in the public signal. α_D is distinguished from α_R, which is the informational discount factor of the regulator constrained by the regulator's recency bias, in order not to impose a priori that it is necessarily optimal for the regulator to disclose its own estimate, as this has to result ex-post from the constrained maximization of its objective function. Therefore, in order to set its optimal communication policy when it does not hold rational expectations, the regulator solves the following maximization problem:

$$\min_{\alpha_D} \text{Var}[\sum_{j=1}^{J-t} \beta^j \exp(0.5(1-\gamma)^2\sigma^2 + (1-\gamma)m_{post,t} + 0.5(1-\gamma)^2\sigma_{post,t}^2 j)^j] \tag{24}$$

s.t.

$$0 \leq \alpha_D \leq \alpha_R,$$

with

$$m_{post,t} = \mathbb{1}_D m_t(\alpha_D) + (1 - \mathbb{1}_D)m_t,$$

and

$$\sigma_{post,t} = \mathbb{1}_D \sigma_t(\alpha_D) + (1 - \mathbb{1}_D)\sigma_t,$$

where $\mathbb{1}_D = 1$ if $\alpha < \alpha_D$ and $\mathbb{1}_D = 0$ if $\alpha \geq \alpha_D$. Indeed, given that estimates (both of the agent and of the regulator) are endogenously formed from the observation of past and current realizations of the logged dividend growth process and consequently are not independent, in the case in which the regulator displays information, the representative investor processes it by substituting it to its own beliefs when the regulator's signal is strictly more precise than his private signal (that is when $\alpha_D > \alpha$) and by ignoring it when it is less precise (that is when $\alpha_D \leq \alpha$). Accordingly, in the case in which the regulator's estimate is less precise than that of the agent (that is, the regulator is more recency-biased than the agent), if the regulator discloses information about the true model's parameter, this does not impact the agent's expectations, as the signal is constrained to be at most as precise as the regulator's estimate. Therefore, it does not impact the investor's optimal decision, and it is optimal for the regulator not to disclose information. On the reverse, when the regulator forms more precise beliefs than the agent, it is able to spread a public signal with higher precision than the agent's private signal. In that case, the agent adopts the estimate disclosed by the regulator when making its decision, as it provides him with more precise information on the true model's parameter, what makes communication policy optimal if and only if it allows to mitigate asset prices excess volatility.

Following the regulator's decision regarding information disclosure, the new price-dividend ratio in each period t writes:

$$\frac{P_t}{D_t} = \sum_{j=1}^{J-t} \beta^j \exp(0.5(1-\gamma)^2\sigma^2 + (1-\gamma)m_{post,t} + 0.5(1-\gamma)^2\sigma^2_{post,t}j)^j. \tag{25}$$

As it depends on $m_{post,t}$ and on $\sigma_{post,t}$, the price-dividend ratio is thus modified by information disclosure as soon as this latter is optimal from the regulator's point of view. In this case, the representative agent's beliefs are affected by the degree of recency bias of the regulator, as this latter restricts the maximal precision of the regulator's public signal.

In the limiting case in which the regulator knows the true parameter in the initial period (that is, it holds rational expectations), $m_{R,t} = d$ and $\sigma_{R,t} = 0$ for $t \geq 1$. In this case, it is optimal for the regulator to disclose information about the true model's parameter and more specifically to disclose its own estimate. Indeed, consequently, agents adopt the regulator's beliefs in the first period and no longer update their beliefs in the following periods, as the precision of their estimate is infinite. The price-dividend ratio following information disclosure writes:

$$\frac{P_t}{D_t} = \sum_{j=1}^{J-t} \beta^j \exp((1-\gamma)dj + 0.5(1-\gamma)^2\sigma^2 j) = \delta_t.$$

For all $1 \leq t \leq J$, it thus reduces to the rational expectations price-dividend ratio, which tends not to vary over time. Therefore, in that extreme case, it is obvious that disclosure of the regulator's own estimate to agents is optimal for mitigating non-fundamental fluctuations in asset prices.

In the general case, it is not possible to derive analytically the optimal α_D of the regulator, as all expressions depend on the sequence of realizations of the data.

However, it is possible to derive analytical results which provide intuition on the mechanism through which information disclosure from the regulator about the true model's parameter can mitigate non-fundamental fluctuations in asset prices and how differences between the case where the regulator is recency-biased and the case where it is not arise.

Thus, the quantity $\frac{\partial m_{post,t}}{\partial y_t}$ measures how the mean belief reacts to the current realization of the data following information disclosure. In the case where agents adopt the estimate disclosed in the public signal in period t (that is they rely on the estimate which is derived based on the parameter α_D), for any $0 \leq \alpha_D < 1$ it writes:

$$\frac{\partial m_{post,t}}{\partial y_t} = \frac{\rho}{\rho \frac{1-\alpha_D^t}{1-\alpha_D} + \alpha_D^t \tau_0} \geq 0. \tag{26}$$

In accordance with intuition, the higher the realization of the logged growth rate of the di- vidends, the higher the mean belief on the mean of the fundamental process. For $\alpha_D = 1$, it writes:

$$\frac{\partial m_{post,t}}{\partial y_t} = \frac{\rho}{\rho t + \tau_0} \geq 0. \tag{27}$$

In both cases, $\frac{\partial^2 m_{post,t}}{\partial y_t \partial \alpha_D}$ is negative. This means that the lower the degree of recency bias, the lower the reaction of the mean belief to new realizations of the data. Intuitively, this stems from the fact that when the degree of recency bias in the estimate disclosed by the regulator is higher, the precision of the information provided by the sample of past data is lower, and thus the current estimate is more sensitive to new data relative to earlier one. Therefore, agents are more likely to overreact to new data realizations, updating their mean belief strongly upwards if the new data is higher than their prior mean belief and strongly downwards if it is lower. Beliefs are thus more volatile. In particular, the variation is clearly higher when $\alpha_D < 1$ in comparison with $\alpha_D = 1$, as $\frac{1-\alpha_D^t}{1-\alpha_D} \leq t$ for any $t > 1$.[11] This is even more obvious when t tends to infinity. Then $\frac{\partial m_{post,t}}{\partial y_t} \rightarrow \frac{\rho}{\frac{\rho}{1-\alpha_D}} = 1 - \alpha_D$ when $\alpha_D < 1$ and $\frac{\partial m_{post,t}}{\partial y_t} \rightarrow 0$ when $\alpha_D = 1$, reflecting the fact that in the absence of recency bias, uncertainty on the model's parameter vanishes in the limit. In addition, it appears that non-linearities in the impact of a marginal decrease in recency bias on the volatility of beliefs arise. This is obvious when the number of observations becomes very high, as $\frac{\partial^2 m_{post,t}}{\partial y_t \partial \alpha_D} \rightarrow -1$ when $\alpha_D < 1$, whereas $\frac{\partial^2 m_{post,t}}{\partial y_t \partial \alpha_D} \rightarrow \frac{-1}{2}$ when $\alpha_D = 1$, implying that a marginal decrease in the recency bias is more efficient for mitigating the volatility in the mean belief when the regulator's recency bias is not negligible.

The impact of the recency bias on the volatility of the dispersion parameter of the posterior distribution σ_t is more ambiguous, except in the limit, where it becomes constant at $\frac{\rho}{1-\alpha}$ when $\alpha_D < 1$ and at 0 when $\alpha_D = 1$ due to convergence to the true parameter. Analytical results thus suggest that when the number of observations is high, it is optimal for the regulator to set $\alpha_D^* = \alpha_R$ if $\alpha_R > \alpha$, as information disclosure about the actual parameter enables to mitigate more non-fundamental volatility in the mean belief, which is the only time-varying factor in the limit. Nevertheless, those results also suggest that such volatility tends to zero only when the regulator is not recency-biased. In addition, unsurprisingly, the gap between post-disclosure volatility in the case with recency bias in the regulator learning process and in the case without recency bias is higher the higher the regulator's degree of recency bias. However, it seems to be more rewarding for the regulator to try to decrease marginally its degree of recency bias when this latter does not tend to zero, as it then decreases more the reaction of the mean belief to the current data realization.

This analytical analysis only enables to provide qualitative results on the efficiency of information disclosure regarding the actual model's parameter for mitigating non-fundamental fluctuations, which justifies to rely on a simulation exercise applied to the recent period in order to assess quantitatively the impact of such policy.

I now illustrate the mechanism of the model with a simulation exercise on the US stock market, starting in the run-up to the most recent bust in the stock market, which followed the initial shock on the subprime loans market. Simulation results then allow to provide some elements of an answer to the maximal recency bias required in order to achieve a significant reduction in

[11] Applying the mean value theorem to the function $f(x) = x^t$ on the interval $[0,1]$, given that $\frac{\partial f}{\partial x} = tx^{t-1} \leq t$, one get $\mid \frac{f(1)-f(\alpha_D)}{1-\alpha_D} \mid \leq t$.

non-fundamental fluctuations. Quantitatively, in accordance with analytical results, it appears that only a very low degree of recency bias in the regulator's learning process enables to achieve a significant decrease in volatility, but that an attempt to decrease marginally the recency bias has more impact on non-fundamental fluctuations when past observations are paid less attention.

5 An illustration on the S&P 500 US stock market in the run-up to the subprime financial crisis

5.1 Simulation results

In order to assess quantitatively the impact of the regulator's recency bias on the efficiency of information disclosure, I now calibrate the model on the US S&P 500 monthly data (Robert Shiller's dataset[12]) in order to replicate some features of the evolution in the price-dividend ratio in the run-up to the 2008-2009 financial crisis. The objective of this small calibration exercise is to focus on one specific period of boom followed by a very strong bust on the US stock market in order to illustrate to what extent information disclosure about the fundamental process could help avoid the occurrence of costly busts such as the one that was observed in the first half of the year 2009 in the US stock market.
Parameters of the real dividend growth process are chosen so as to match those in the data over the period January 1960-June 2009 (Table 1). Indeed, as shown by Pesaran et al. (2007), a structural break in the dividend process occurred in the year 1960. June 2009 corresponds to the end of our period of study, just after the climax of the bust was reached. Subsequent data is likely to present structural breaks and the learning process agents relied on in the former period is likely to have suffered strong changes, both in terms of underlying parameters and learning mechanism. The monthly discount rate β is set at 0.998, based on prior literature in which it is standard to assume a quarterly discount rate of 0.99. The monthly informational discount rate α is set to 0.92, relying on empirical studies which have evaluated it in distinct markets (see in particular (Agarwal et al., 2013)). The constant relative risk aversion parameter for risk-averse investors $0 < \gamma$ (which is also the inverse of the intertemporal elasticity of substitution) is constrained to be inferior to 1 in order not to produce counter-intuitive features that would contradict the

[12]Monthly data on the US S&P 500 stock market is retrieved from Shiller's dataset available online.

data. Indeed, if $\gamma > 1$, then the wealth effect dominates the substitution effect, meaning that when the discounted expected present value of future real dividends streams is expected to be higher, demand for stocks decreases whereas consumption increases, leading to a decrease in stock prices. Nevertheless, as empirical assessments of risk aversion coefficients usually show them to be relatively high, γ is chosen so as to be the closest possible to 1 while still generating a consistent order of magnitude for the rational expectations value of the price-dividend ratio.
J is chosen so as to be high enough not to impact prices over the simulation period but such that simulations can still be computationally possible.
As the prior distribution $d \sim N(m_0, \sigma_0)$ sums up the initial information the agent has on the estimated parameter, the parameters m_0 and σ_0 are derived from pre-sample dividend realizations. As a structural break occurred in the dividend process around the year 1960, the prior parameters are retrieved from dividend realizations over the period 1960-2003 in a way which is consistent with the in-sample agents' learning process: m_0 is the weighted average of past observations relying on the informational discount rate α,[13] with the realization in January 1960 being discounted the most and the realization in December 2002 not being discounted, and σ_0^2 is the inverse of the discounted sum of the precision of the information provided by each data point, that is of the precision of the fundamental process ρ. Finally, the model is fed with the exact same dividends realizations as those observed in the data in the period of interest when deriving the model-implied price-dividend ratio. The subjective expectations model performs

Parameter	Calibrated value
d	0.0011
σ	0.0056
β	0.998
γ	0.9
α	0.92
m_0	-0.0013
σ_0	0.0016
J	20 000

Table 1: Calibration

strikingly much better than the rational expectations benchmark. First and second order short-

[13]Note that agents' prior belief, based on pre-sample dividend information, is such that they undervalue the actual parameter (Table 2).

term moments of the subjective expectations model-implied price-dividend ratio over the period of interest replicate rather well those in the data, even if they have not been directly targeted in the calibration strategy (Table 2). The mean of the subjective expectations price-dividend ratio is significantly above its rational expectations value, which reflects the boom episode on the US stock market in the run-up to the recent financial crisis. The subjective expectations price-dividend ratio standard error is much closer to that in the data than the rational expectations one, which is negligible (but non-null due to finite time). The model generates though slightly too much volatility in comparison with what is observed in the data. However, despite the simplicity of the model and the fact that the only source of variation in the price-dividend ratio in the model is the variation in beliefs regarding one parameter only, some qualitative features of the data are also replicated to a relatively good extent (Figure 1).

First, the model replicates the boom period in the US stock market in the aftermath of the dot-com bubble bust, with the price-dividend ratio increasing significantly and remaining persistently above its rational expectations value (what can be identified as a bubble). Second, the model replicates the deep decrease in the price-dividend ratio, reaching values well below the rational expectations value (the bust). The bust in the price-dividend ratio results from the transmission of a strong negative dividend shock (which may reflect exogenous deteriorating financial and economic conditions in the context of the subprime crisis) to beliefs on the actual fundamental process due to persistent parameter uncertainty, and then to demand for the risky asset. As agents are risk averse, this makes stock prices collapse. The extent of the bust is all the more so strong that the actual parameter of the fundamental process was overvalued in the previous periods due to a series of positive fundamental shocks. Indeed, the strong negative shock thus induces a higher reassessment of the parameter estimate. The simulation exercise thus makes obvious how the alternation of phases of overvaluation and undervaluation of stocks' fundamental value generates significant booms and busts episodes in the stock market. Finally, these results suggest that time-varying backward-looking expectations played a non-negligible role in explaining stock prices fluctuations over the recent boom and bust period, all the more so that the model is able to replicate additional features of the data.

Thus, firstly, the model replicates well the strong autocorrelation in the price-dividend ratio behavior over the period (Table 2). Secondly, the dynamics of the realized monthly stock returns are more consistent with those observed empirically than those generated by the rational ex-

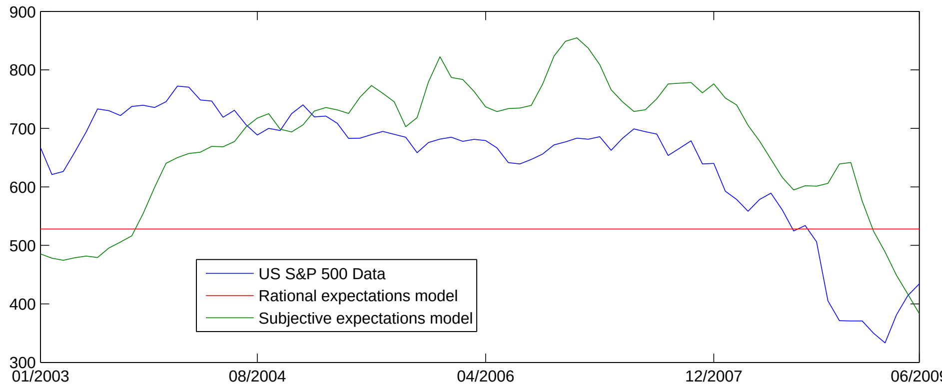

Figure 1: 2003-2009 Monthly price-dividend ratio

pectations benchmark model (Figure 2). Third, the subjective expectations model generates a positive correlation between expected stock returns, as observed in survey data on expected returns (Table 2),[14] even though it tends to overestimate it.

	RE model	SE model	Data
P/D ratio Mean	528.0844	674.3400	640.7275
P/D ratio Standard deviation	$4.5769*10^{-13}$	113.5691	108.9037
Autocorrelation in the P/D ratio	1 (0.0000)	0.9747 (0.0000)	0.9777 (0.0000)
Correlation between one-year ahead expected returns (CFO survey) and the price-dividend ratio			
Monthly	NaN	0.9837 (0.0000)	0.7420 (0.0000)
Quarterly	NaN	0.9814 (0.0000)	0.6743 (0.0002)

Table 2: Simulation results

Finally, in order to make clear what drives the dynamics of asset prices in the model, Figure 3 displays the joint dynamics of model-implied beliefs and the price-dividend ratio. It shows how the mean belief m_t regarding the value of the true parameter d (left-hand scale) fluctuates around d, with first a sustained period of optimism –in which $m_t > d$– and second a sudden peak of pessimism –in which m_t goes far below d–, which drives the dynamics of the price-dividend ratio (right-hand scale). The uncertainty parameter σ_t being constant in the simulation

[14]Model-implied one-year ahead expected returns are simply annualized monthly returns. The CFO survey being published quarterly and the model time period being a month, for comparability issues, I linearly interpolate quarterly expected returns in the CFO survey in order to get monthly data. In order to get an additional result which is independent of interpolation methods, I compare quarterly data as well, by taking model-implied annualized monthly expected returns of the last month of each quarter as a proxy for each quarter one-year ahead expected returns.

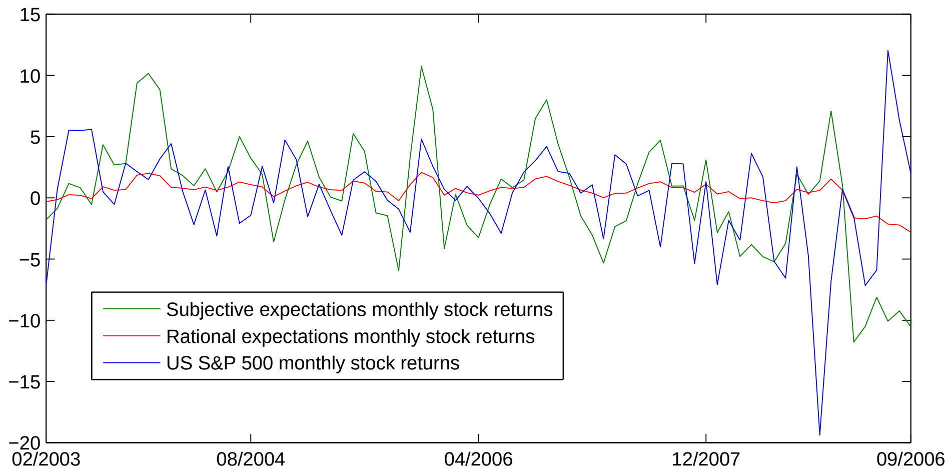

Figure 2: 2003-2009 Monthly net returns on stocks (%)

results,[15] it is obvious that the dynamics of the mean belief parameter directly translate into the dynamics of the price-dividend ratio. Consequently, this enables us to assess quantitatively un-

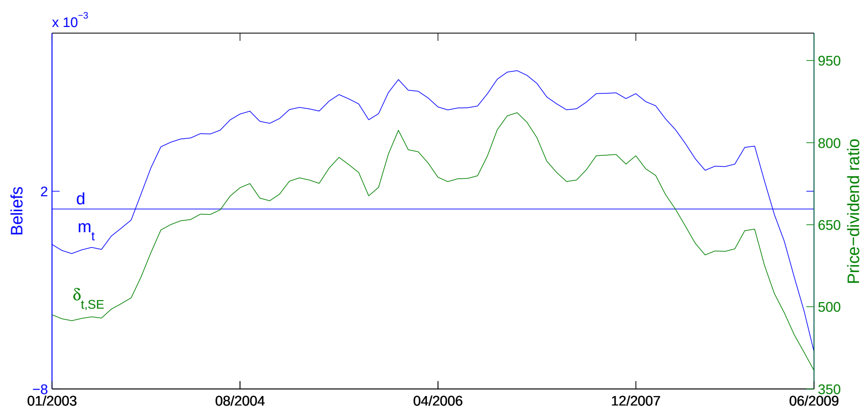

Figure 3: Joint dynamics of beliefs and the price-dividend ratio

der which condition on the regulator's precision of information, the alternation of booms and

[15]This is due to informational discounting with prior information gathering a significant number of past pre-sample observations, which implies that even though more data is accumulated by agents over time, precision of information is not improved as the earliest data is allocated a zero weight. For the given model's parameter values, the constant precision of agents' information on the unknown parameter is equal to $3.9 * 10^5$ (that is the uncertainty on this parameter is equal to $2.5 * 10^{-6}$). Therefore, the results show that even for a very low value of uncertainty, beliefs' updating still generates significant volatility in the price-dividend ratio.

busts episodes in the price-dividend ratio can be significantly mitigated following information disclosure.

5.2 The impact of information disclosure

I now assess quantitatively under which condition on the precision of the public signal on the parameter of the actual fundamental process – which depends on α_D – information disclosure can be a relevant solution for mitigating volatility in asset prices through a simple stylized counterfactual simulation exercise. I consider only cases in which $\alpha_R \geq \alpha$ (that is the regulator is less recency-biased than economic agents). As explained above, when $\alpha_R < \alpha$, the private signal of the representative agent is necessarily more precise than that of the regulator (which is not independent), and it is optimal for the investor just to ignore it. The price-dividend ratio thus remains unchanged.
Figure 4 presents the evolution in the price-dividend ratio following information disclosure for distinct levels of the regulator's informational discount rate α_R (and thus for distinct degrees of the regulator's recency bias or alternatively for distinct degrees of precision of its information on the actual parameter). As $\alpha_D \leq \alpha_R$ (meaning that the informational discount factor used in the derivation of the public signal is not restricted a priori to be equal to the regulator's constrained informational discount factor but cannot logically be higher), results obtained by making α_R vary hold for α_D, which is the decision variable of the regulator when setting its information disclosure policy.

At first glance, it is striking that there are strong differences in the impact of information disclosure on the volatility of the price-dividend ratio, depending on the regulator's degree of recency bias. First, volatility in the price-dividend ratio tends to zero only when the regulator is not recency-biased, confirming analytical results. Even for a very small informational discount rate of $\alpha_R = 0.98$, the post-information disclosure price-dividend ratio remains significantly away from its rational expectations 'fundamental' value. However, decreasing the recency-bias always decreases the volatility, what suggests that $\alpha_D^* = \alpha_R$. Second, it appears that the impact of similar decreases in the regulator's degree of recency bias does not impact the price-dividend ratio in the same extent whatever the level of the regulator's informational discount rate, reflecting the existence of non-linearities in the impact of an increase in the precision of the regulator's

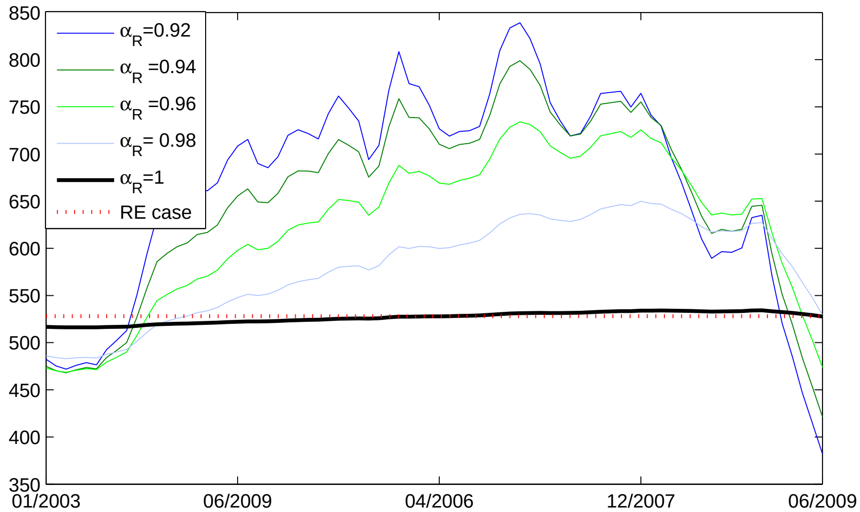

Figure 4: The impact of information disclosure on the price-dividend ratio for various degrees of the regulator's recency bias

information on the price-dividend ratio volatility. To complement this finding, the following chart presents the variance of the subjective expectations price-dividend ratio and the mean of the squared distance of the subjective expectations price-dividend ratio to its rational expectations value as functions of the regulator's recency bias.

First, it appears that a lower degree of recency bias monotonically reduces the variance of the subjective expectations price-dividend ratio and its average squared distance to the rational expectations value but both variables tend to zero only when the precision of the regulator's information under subjective expectations is maximal. Second, a marginal increase in the precision of the regulator's information (due to lower recency bias) seems to be more efficient in mitigating non-fundamental fluctuations in the price-dividend ratio when the regulator's informational discount rate is higher, to the striking exception of the case where it is already very close to 1.

Finally, I assess under which condition on the degree of the regulator's recency bias distinct objectives in terms of mitigating the volatility in the price-dividend ratio and bringing it closer to its rational expectations value could be achieved.[16] The simulation results are displayed in Table 3. It appears that very small degrees of recency bias in the regulator learning process in

[16] As a comparison, when there is no information disclosure, the standard error of the price-dividend ratio over time is equal to 16.84% of its mean, and the squared root of the mean squared distance of the price-dividend ratio to its rational expectations value is equal to 34.98% of the rational expectations value.

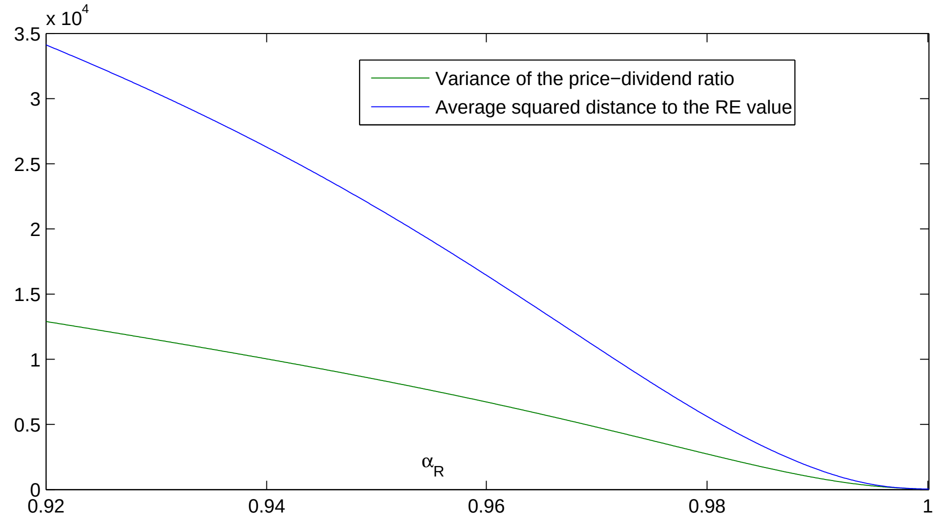

Figure 5: Statistical features of the price-dividend ratio after information disclosure depending on the regulator's degree of recency bias α_R

Objective	Minimal α_R (order -4)
Standard error of P/D: <5% of the mean	0.9908
Standard error of P/D: <10% of the mean	0.9762
Squared root of the mean squared distance of P/D to its RE value: <5% of the RE value	0.9933
Squared root of the mean squared distance of P/D to its RE value: <10% of the RE value	0.9864
Squared root of the mean squared distance of P/D to its RE value: <20% of the RE value	0.9695

Table 3: Minimal degree of the regulator's recency bias required in order to achieve distinct objectives

comparison with the agent's recency bias are required so as to achieve significant decrease in the price-dividend ratio volatility and in the average distance of the price-dividend ratio to its rational expectations value. Therefore, if information disclosure about the actual fundamental process seems to be a relevant tool whenever the regulator is not recency-biased, as soon as it is itself recency-biased, this raises serious concerns on its ability to significantly mitigate non-fundamental fluctuations in asset prices. These results thus suggest that, in order to make information disclosure a useful tool in mitigating inefficient fluctuations in asset prices, more attention has to be paid to longer-span historical series of data, as recommended by Haldane (2009) or Reinhart and Rogoff (2009).

6 Conclusion

A parsimonious standard consumption-based asset pricing model in which agents learn the location parameter of the dividend growth process through recency-biased Bayesian inference, providing microfoundations to investors' decision without implying any restrictive assumption on agents' knowledge of the pricing function, enables to derive a closed-form solution for stock price. This makes obvious how the latter depends on investors' expectations and how this triggers fluctuations in the price-dividend ratio and thus generates the potential for non-fundamental bubbles. The specificity of the model is that the extent and the persistence of these fluctuations over time are due to the representative investor's recency bias, relying on growing empirical evidence.

Even with a small degree of parameter uncertainty, the model proves able to replicate several features of the US stock market in the run-up to the subprime crisis: the price-dividend ratio displays significant volatility over time and evolves according to surprise effects –thus displaying a steep decrease in the beginning of 2009–, it is strongly autocorrelated, and positively correlated with expected future returns. The model also replicates qualitative features of the dynamics of stock returns.

Modelling the dynamics of subjective expectations in an otherwise standard asset prices model thus leads to new predictions relative to those of rational expectations models. It enables to predict that, unsurprisingly, expectations-driven booms and busts do arise. This paves the way for information disclosure from the regulator about the actual parameter of the fundamental process in order to warn market participants against possible over or undervaluation of assets. Nevertheless, information disclosure about the actual parameter significantly mitigates expectations-driven fluctuations only when the regulator's recency bias tends to zero.

Those results suggest that recency bias in the regulator's learning process may deprive financial regulation authorities of a tool that could otherwise prove useful in mitigating excess volatility in asset prices. To this aim, it matters that persistent attention is paid not only to recent data but also to earlier historical ones. This would allow to better identify unusual behavior in asset prices and other macro-financial variables relative to their historical behavior and would prevent mixing-up transitory recent trends with permanent structural evolution, which has potentially disastrous consequences for financial stability.

References

Adam, K., Beutel, J., and Marcet, A. (2015a). Stock Price Booms and Expected Capital Gains. *Mimeo*.

Adam, K., Beutel, J., Marcet, A., and Merkel, S. (2014). Can a Financial Transaction Tax Prevent Stock Price Booms ? *Mimeo*.

Adam, K., Marcet, A., and Nicolini, J. (2015b). Stock Market Volatility and Learning. *Journal of Finance*, forthcoming.

Agarwal, S., Driscoll, J., Gabaix, X., and Laibson, D. (2013). Learning in the Credit Card Market. *NBER Working Paper*.

Bansal, R. and Shaliastovich, I. (2010). Confidence Risk and Asset Prices. *The American Economic Review*, 100(2):537:541.

Bloxham, P., Kent, C., and Robson, M. (2010). Asset Prices, Credit Growth, Monetary and Other Policies: An Australian Case Study. *Research Discussion Paper, Reserve Bank of Australia*, 06.

Cheung, Y. and Friedman, D. (1997). Individual Learning in Normal Form Games: Some Laboratory Results. *Games and Economic Behavior*, 19(1):46-76.

Cooper, J. and Kovacic, W. (2012). Behavioral Economics: Implications for Regulatory Behavior. *Journal of Regulatory Economics*, 41:41-58.

de Bondt, W. and Thaler, R. (1985). Does the Stock Market Overreact ? *Journal of Finance*, 40:793-805.

de Bondt, W. and Thaler, R. (1990). Do Security Analysts Overreact ? *The American Economic Review*, 80:52-57.

Erev, I. and Haruvy, E. (2013). *Learning and the Economics of Small Decisions*, volume 2. Roth, A.E. Kagel, J. Eds.

Fudenberg, D. and Levine, D. (2014). Recency, Consistent Learning and Nash Equilibrium. *Proceedings of the National Academy of Sciences*, 111: 10826-10829.

Gallagher, J. (2014). Learning from an Infrequent Event: Evidence from Flood Insurance Take-up in the United States. *American Economic Journal: Applied Economics*, 6(3): 206-233.

Haldane, A. (2009). Why Banks Failed the Stress Test ? In *Speech given at the Marcus-Evans Conference on Stress-Testing*.

Koulovatianos, K. and Wieland, W. (2011). Asset Pricing Under Rational Learning About Rare Disaster Risk. *CEPR Discussion Working Paper*, 8514.

LeRoy, S. and Parke, W. (1992). Stock price volatility: Tests based on the geometric random walk. *The American Economic Review*, 981-992.

Lovell, C. (1986). Tests of the Rational Expectations Hypothesis. *The American Economic Review*, Vol.76(1):110-124.

Lucas, R. (1978). Asset Prices in an Exchange Economy. *Econometrica*, 46-6:1429-1445.

Malmendier, U. and Nagel, S. (2011). Depression Babies: Do Macroeconomic Experiences Affect Risk Taking ? *The Quarterly Journal of Economics*, 126(1): 373-416.

Mishkin, F. (2011). How Should Central Banks Respond to Asset Bubbles ? *Reserve Bank of Australia Bulletin*, 59.

Nakov, A. and Nuno, G. (2015). Learning from Experience in the Stock Market. *Journal of Economic Dynamics and Control*, 52:224-239.

Pesaran, H., Pettenuzzo, D., and Timmermann, A. (2007). Learning, Structural Instability, and Present Value Calculations. *Econometric Reviews*, 26: 253-288.

Reinhart, C. and Rogoff, K. (2009). *This Time is Different: Eight Centuries of Financial Folly*. Princeton University Press.

Timmerman, A. (1993). How Learning in Financial Markets Generates Excess Volatility and Predictability in Stock Prices ? *The Quarterly Journal of Economics*, 108(4):1135-1145.

Timmermann, A. (1996). Excess Volatility and Predictability of Stock Prices in Autoregressive Dividend Models with Learning. *Review of Economic Studies*, 63(4):523-557.

Williams, J. (2014). Financial Stability and Monetary Policy: Happy Mariage or Untenable Union? *FRBSF Economic Letters*, 17.

Winkler, F. (2015). Learning in the Stock Market and Credit Frictions. *Mimeo*.

Woodford, M. (2012). Inflation Targeting and Financial Stability. *NBER Working Paper*, 17967.

A Proof of Proposition 1

The first Euler equation (with respect to quantity of stocks S_t) after market clearing writes:

$$D_t^{-\gamma} = \beta E_t[D_{t+1}^{-\gamma}(\frac{P_{t+1} + D_{t+1}}{P_t})],$$

for $1 \leq t \leq J-1$. Isolating P_t on the left hand side yields:

$$P_t = \beta E_t[(\frac{D_{t+1}}{D_t})^{-\gamma}(P_{t+1} + D_{t+1})].$$

Substituting P_{t+1} by its expression in the iterated forward version of the previous equation leads to:

$$P_t = \beta E_t[(\frac{D_{t+1}}{D_t})^{-\gamma}(\beta E_{t+1}[(\frac{D_{t+2}}{D_{t+1}})^{-\gamma}(P_{t+2} + D_{t+2})]) + (\frac{D_{t+1}}{D_t})^{-\gamma} D_{t+1}]$$

Applying the law of iterated expectations with nested conditioning sets ($E_t[E_{t+1}(X)] = E_t[X]$) and iterating forward again yields:

$$P_t = E_t[\beta^{J-t}(\frac{D_J}{D_t})^{-\gamma} P_J] + E_t[\beta(\frac{D_{t+1}}{D_t})^{-\gamma} D_{t+1} + \beta^2(\frac{D_{t+2}}{D_t})^{-\gamma} D_{t+2} + ... + \beta^{J-t}((\frac{D_J}{D_t})^{-\gamma} D_J))].$$

In the last period J, under the non-bequest assumption according to which all remaining wealth at the beginning of period J is consumed in J, stocks are no longer traded and thus $P_J = 0$. Finally,

$$P_t = E_t \sum_{j=1}^{J-t} \beta^j (\frac{D_{t+j}}{D_t})^{1-\gamma} D_t.$$

Thus,

$$E_t[(\frac{D_{t+j}}{D_{t+j-1}} * \frac{D_{t+j-1}}{D_{t+j-2}} * ... * \frac{D_{t+1}}{D_t})^{1-\gamma}] = E_t[E_{t+1}[(\frac{D_{t+1}}{D_t})^{1-\gamma} * ... * E_{t+j-1}[(\frac{D_{t+j}}{D_{t+j-1}})^{1-\gamma}]]].$$

Hence, when d and σ are known:

$$E_{t+j-1}[(\frac{D_{t+j}}{D_{t+j-1}})^{1-\gamma}] = E_{t+j-2}[(\frac{D_{t+j-1}}{D_{t+j-2}})^{1-\gamma}] = E_t[(\frac{D_{t+1}}{D_t})^{1-\gamma}] = \exp(d(1-\gamma) + \frac{(1-\gamma)^2\sigma^2}{2}) = \theta.$$

Therefore,

$$E_t[(\frac{D_{t+j}}{D_t})^{1-\gamma}] = \theta^j,$$

and

$$P_t = D_t \sum_{j=1}^{J-t} \beta^j \theta^j.$$

It is the sum of the J-t first terms of a geometric sequence with common ratio $\beta\theta$ and first term $\beta\theta$. Therefore,

$$P_t = \frac{\beta\theta - (\beta\theta)^{J-t+1}}{1-\beta\theta} D_t.$$

B Inconsistent features of the rational expectations model

B.1 Constant versus highly volatile price-dividend ratio

The following chart displays the US S&P 500 monthly price-dividend ratio over the recent period and the rational expectations model implied one, which is (roughly) constant (the parameters values used here are the same as those presented in the simulation exercise in Section 5.)

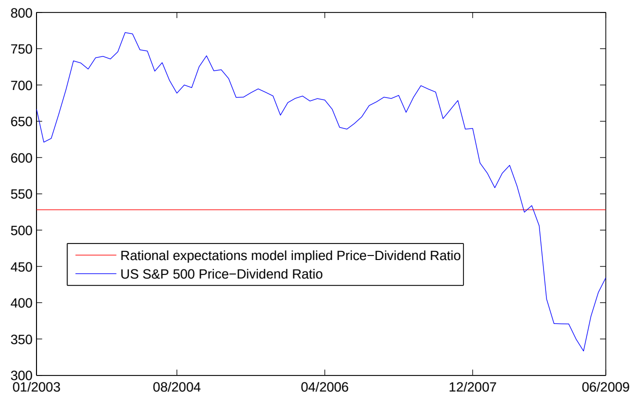

Monthly 2003-2009 US S&P 500 price-dividend ratio versus the rational expectations benchmark implied one

B.2 Volatile stock returns

The following chart presents realized monthly returns on US S&P 500 stocks. They display very low volatility, what strikingly contradicts the data.

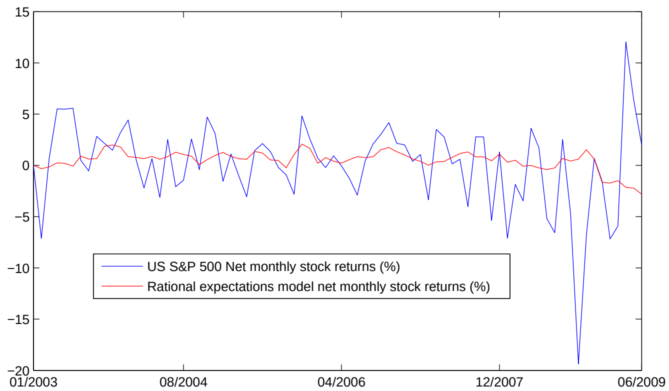

Monthly stock market returns: data versus rational expectations model

B.3 Positive correlation between one-year ahead time-varying expected stock returns (CFO Survey) and the current price-dividend ratio

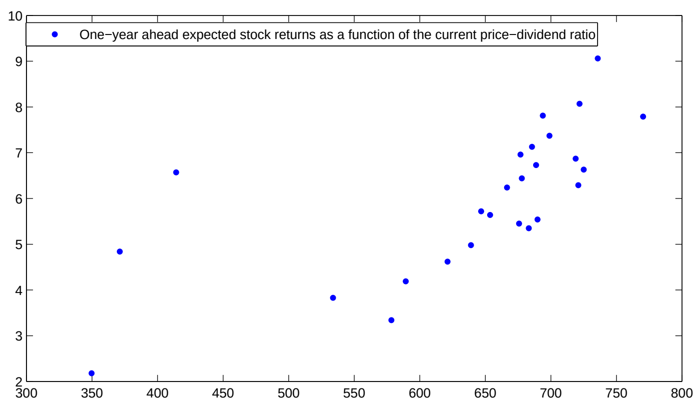

Expected stock returns as a function of the price-dividend ratio

C Derivation of the posterior distribution under recency-biased learning

Under recency-biased learning, in period t, the precision of the prior distribution is discounted by α^t. Thus, the normal prior distribution $P(d \mid I_0, \sigma)$ writes:

$$P(d \mid I_0, \sigma) = \sqrt{\frac{\alpha^t \tau_0}{2\pi}} \exp(-\frac{1}{2}\alpha^t \tau_0 (d - m_0)^2).$$

Similarly, the precision of the realization of the data y_k in period $k < t$ is discounted by α^k in the joint likelihood:

$$L(y^t \mid d, \sigma) = \prod_{k=0}^{t-1} \sqrt{\frac{\alpha^k \tau}{2\pi}} \exp(-\frac{1}{2}\alpha^k \tau (y_{t-k} - d)^2).$$

Therefore, the posterior distribution $P(d \mid I_t, \sigma)$ is:

$$P(d \mid I_t, \sigma) \propto \sqrt{\frac{\alpha^t \tau_0}{2\pi}} \prod_{k=0}^{t-1} (\sqrt{\frac{\alpha^k \tau}{2\pi}}) \exp(-\frac{1}{2}\alpha^t \tau_0 (d - m_0)^2) \exp(-\frac{1}{2}\tau (\sum_{k=0}^{t-1} \alpha^k ((y_{t-k} - \bar{y_W})^2 + \sum_{k=0}^{t-1} \alpha^k (\bar{y_W} - d)^2)),$$

with $\bar{y_W} = \frac{\sum_{k=0}^{t-1} \alpha^k y_{t-k}}{\sum_{k=0}^{t-1} \alpha^k}$ the weighted average of past observations. This yields:

$$P(d \mid I_t, \sigma) \propto \exp(-\frac{1}{2}(\alpha^t \tau_0 (d - m_0)^2 + \tau \sum_{k=0}^{t-1} \alpha^k (d - \bar{y_W})^2))$$

Completing the square thus leads to:

$$P(d \mid I_t, \sigma) \propto \exp(-\frac{1}{2}(\alpha^t \tau_0 + \tau \sum_{k=0}^{t-1} \alpha^k)(d - \frac{\alpha^t \tau_0 m_0 + \tau \sum_{k=0}^{t-1} \alpha^k \bar{y_W}}{\alpha^t \tau_0 + \tau \sum_{k=0}^{t-1} \alpha^k})^2).$$

Therefore, $P(d \mid I_t, \sigma) \sim N(\frac{\alpha^t \tau_0 m_0 + \tau \sum_{k=0}^{t-1} \alpha^k \bar{y_W}}{\alpha^t \tau_0 + \tau \sum_{k=0}^{t-1} \alpha^k}, \frac{1}{\sqrt{\alpha^t \tau_0 + \tau \sum_{k=0}^{t-1} \alpha^k}})$.

D Derivation of equation (14)

Starting from the Euler equation with respect to stocks under the subjective probability measure, isolating P_t on the left-hand side and iterating forward as in Proof of Proposition 1 yields:

$$P_t = \sum_{j=1}^{J-t} \beta^j D_t E_t^P[(\frac{D_{t+1}}{D_t})^{1-\gamma} * E_{t+1}^P[(\frac{D_{t+2}}{D_{t+1}})^{1-\gamma} * ... * E_{t+j-1}^P[(\frac{D_{t+j}}{D_{t+j-1}})^{1-\gamma}]]],$$

$$= \sum_{j=1}^{J-t} \beta^j D_t E_t^P[E_{t+1}^P[E_{t+2}^P[...E_{t+j-1}^P[(\frac{D_{t+1}}{D_t})^{1-\gamma} * (\frac{D_{t+2}}{D_{t+1}})^{1-\gamma} * ... * (\frac{D_{t+j}}{D_{t+j-1}})^{1-\gamma}]]]],$$

for $1 \leq t \leq J-1$. Under the assumption of adaptive Bayesian learning, that is when agents take into account the uncertainty on their estimates but not the probability distributions of future beliefs (that is, agents do not take into account the fact that their current beliefs may change following future realizations of the dividends growth rate), the law of iterated expectations can still be applied, as $E_t[E_{t+1}(X)] = E_t[X]$. Eventually, this yields:

$$P_t = \sum_{j=1}^{J-t} \beta^j D_t E_t^P[(\frac{D_{t+j}}{D_t})^{1-\gamma}].$$

E Comparative statics

As $P_t = D_t \sum_{j=1}^{J-t} \beta^j \exp((1-\gamma)m_t j + 0.5(1-\gamma)^2 j^2(\sigma_t^2 + \sigma^2))$, the sign of its derivative with respect to any variable is the same as the sign of the derivative of the term in the sum with respect to that variable. For convenience, I now write $\beta^j \exp((1-\gamma)m_t j + 0.5(1-\gamma)^2 j^2(\sigma_t^2 + \sigma^2)) = x_t$.

$$\frac{\partial x_t}{\partial \beta} = j\beta^{j-1} \exp((1-\gamma)m_t j + 0.5(1-\gamma)^2 j^2(\sigma_t^2 + \sigma^2)) > 0.$$

When the discount rate is higher –that is preference for current consumption is lower– the demand for stock prices increases and thus the equilibrium stock price increases.

$$\frac{\partial x_t}{\partial \gamma} = \beta^j j \exp((1-\gamma)m_t j + 0.5(1-\gamma)^2 j^2(\sigma_t^2 + \sigma^2))(-m_t - (1-\gamma)(\sigma^2 + \sigma_t^2)j).$$

$$\frac{\partial x_t}{\partial \gamma} < 0 \Leftrightarrow m_t > (\gamma - 1)(\sigma^2 + \sigma_t^2)j.$$

The sign of the derivative of the stock price with respect to the relative risk aversion coefficient depends on the hyperparameters m_t and σ_t and thus on the data realizations. When relative risk aversion increases, the demand for stocks decreases if and only if the expected mean of the logged growth rate of the payoff on stocks is high enough.

$$\frac{\partial x_t}{\partial \sigma} = \beta^j j^2 \exp((1-\gamma)m_t j + 0.5(1-\gamma)^2 j^2(\sigma_t^2 + \sigma^2))(1-\gamma)^2\sigma > 0.$$

The demand for stocks increases when the (known) variance of the logged growth rate of the payoff on stocks increases.

$$\frac{\partial x_t}{\partial \sigma_t} = \beta^j j^2 \exp((1-\gamma)m_t j + 0.5(1-\gamma)^2 j^2(\sigma_t^2 + \sigma^2))(1-\gamma)^2\sigma_t > 0.$$

Similarly, the demand for stocks increases when the uncertainty on the expected mean of the logged growth rate of the payoff on stocks increases.

$$\frac{\partial x_t}{\partial m_t} = \beta^j j \exp((1-\gamma)m_t j + 0.5(1-\gamma)^2 j^2(\sigma_t^2 + \sigma^2))(1-\gamma).$$

$\frac{\partial x_t}{\partial m_t} > 0 \Leftrightarrow \gamma < 1$. When the substitution effect dominates the wealth effect, demand for stocks increases when the expected mean of the logged growth rate of the payoff on stocks increases.

F Expected Returns

Conditionally on information up to time t and the value of the dividend process parameters d and σ, expected returns write:

$$E_t[\frac{P_{t+1} + D_{t+1}}{P_t} \mid d, \sigma] = \frac{1}{P_t}(E_t[P_{t+1} \mid d, \sigma] + D_t E_t[\frac{D_{t+1}}{D_t} \mid d, \sigma]).$$

In particular,

$$E_t[P_{t+1} \mid d, \sigma] = E_t[D_{t+1} \sum_{j=1}^{J-t-1} \beta^j E_{t+1}[(\frac{D_{t+1+j}}{D_{t+1}})^{1-\gamma} \mid d, \sigma] \mid d, \sigma]$$

$$= D_t \sum_{j=1}^{J-t-1} \beta^j \exp(d((1-\gamma)j+1) + 0.5((1-\gamma)^2 j + 1)\sigma^2).$$

Let's now assume that the parameter d is no longer known and is believed to follow a normal distribution with parameters m_t and σ_t in period t. In this case,

$$E_t[P_{t+1} \mid \sigma] = D_t \sum_{j=1}^{J-t-1} \beta^j \exp(m_t((1-\gamma)j+1) + 0.5((1-\gamma)j+1)^2\sigma_t^2 + 0.5((1-\gamma)^2 j + 1)\sigma^2),$$

assuming that the law of iterated expectations still applies, because agents are adaptive Bayesian learners.

Eventually,

$$E_t[\frac{P_{t+1}+D_{t+1}}{P_t} \mid \sigma] = \frac{D_t}{P_t}(\sum_{j=1}^{J-t-1}[\beta^j \exp(m_t((1-\gamma)j+1) + 0.5((1-\gamma)j+1)^2\sigma_t^2$$
$$+0.5((1-\gamma)^2 j+1)\sigma^2)] + \exp(m_t + 0.5(\sigma_t^2 + \sigma^2))).$$